

3FabRec: Fast Few-shot Face alignment by Reconstruction

Björn Browatzki and Christian Wallraven*
Dept. of Artificial Intelligence, Korea University, Seoul
browatbn@korea.ac.kr, wallraven@korea.ac.kr




Figure 1: The framework of 3FabRec consisting of: (leftmost box) a first, unsupervised training stage that trains a low-dimensional latent space via an adversarial (generative) autoencoder on a large dataset of unlabeled faces, (middle box) subsequent supervised training with few annotated faces. The rightmost box shows results from testing the framework trained on *only the 10 images of the middle box* with original faces (top row), reconstructed faces via the autoencoder (middle row), and confidence heatmaps (bottom row).

Abstract

Current supervised methods for facial landmark detection require a large amount of training data and may suffer from overfitting to specific datasets due to the massive number of parameters. We introduce a semi-supervised method in which the crucial idea is to first generate implicit face knowledge from the large amounts of unlabeled images of faces available today. In a first, completely unsupervised stage, we train an adversarial autoencoder to reconstruct faces via a low-dimensional face embedding. In a second, supervised stage, we interleave the decoder with transfer layers to retask the generation of color images to the prediction of landmark heatmaps. Our framework (3FabRec) achieves state-of-the-art performance on several common benchmarks and, most importantly, is able to maintain impressive accuracy on extremely small training sets down to as few as 10 images. As the interleaved layers only add a low amount of parameters to the decoder, inference runs at several hundred FPS on a GPU.

1. Introduction

Accurate and robust localization of facial landmarks is a critical step in many existing face processing applications, including tracking, expression analysis, and face identification. Unique localization of such landmarks is severely affected by occlusions, partial face visibility, large pose variations, uneven illumination, or large, non-rigid deformations during more extreme facial expressions [41, 44]. These challenges have to be overcome in order to achieve a low landmark localization error, implying high robustness to appearance changes in faces while guaranteeing high localization accuracy for each landmark.

The recent advances in deep learning techniques [3] coupled with the availability of large, annotated databases have allowed steady progress with localization accuracy on a typical benchmark increasing by 100% (from [51] to [49] - see below for more related work). Most approaches use a combination of highly-tuned, supervised learning schemes in order to achieve this performance and almost always

are specifically optimized on the particular datasets that are tested, increasing the potential of overfitting to that dataset [7]. Similarly, it has been shown that annotations in datasets can be imprecise and inconsistent (e.g., [14]).

Given that in addition to the existing annotated facial landmark datasets, there is an even larger number of datasets available for other tasks (face detection, face identification, facial expression analysis, etc.), it should be possible to leverage the implicit knowledge about face shape contained in this pool to both ensure better generalizability across datasets and easier and faster, few-shot training of landmark localization. Here, we present such a framework that is based on a two-stage architecture (3FabRec, see Figs.1,2): the key to the approach lies in the first, *unsupervised* stage, in which an (generative) adversarial auto-encoder [30] is trained on a large dataset of faces that yields a low-dimensional embedding capturing "face knowledge" [46] from which it is able to reconstruct face images across a wide variety of appearances. With this embedding, the second, *supervised* stage then trains the landmark localization task on annotated datasets, in which the generator is retasked to predict the locations of a set of landmarks by generating probabilistic heatmaps [5]. This two-stage approach is a special case of semi-supervised learning [24,59] and has been successful in other domains, including general network training [20], text classification [22] and translation [13], and visual image classification [54].

In the current study, we show that our method is able to achieve state-of-the-art results running at >300 FPS on the standard benchmark datasets. Most importantly, it yields impressive localization performance already with *a few percent of the training data* - beating the leading scores in all cases and setting new standards for landmark localization from as few as *10 images*. The latter result demonstrates that landmark knowledge has, indeed, been implicitly captured by the unsupervised pre-training. Additionally, the reconstructed autoencoder images are able to "explain away" extraneous factors (such as occlusions or make-up), yielding a best-fitting face shape for accurate localization and adding to the explainability of the framework.

Source code is available at <https://github.com/browatbn2/3FabRec>.

2. Related Work

Before the advent of deep learning methods, explicitly-parametrized landmark models such as active shape [9], active appearance [8] or cascade regression models [17, 51] provided the state-of-the-art in facial landmark detection. Current models using deep convolutional neural networks, however, quickly became the best-performing approaches, starting with deep alignment networks [42], fully-convolutional networks [27], coordinate regression mod-

els [28, 45], or multi-task learners [36], with the deep networks being able to capture the pixel-to-landmark correlations across face appearance variations.

The recent related work in the context of our approach can be structured into supervised and semi-supervised approaches (for a recent, interesting unsupervised method - at lower performance levels - see [43]).

2.1. Supervised methods

Several recent, well-performing supervised methods are based on heatmap regression, in which a deep network will infer a probabilistic heatmap for each of the facial landmarks with its corresponding maximum encoding the most likely location of that landmark [5, 12, 27] - an approach we also follow here. In order to provide additional geometric constraints, extensions use an active-appearance-based model-fitting step based on PCA [31], explicit encoding of geometric information from the face boundary [49], or additional weighting from occlusion probabilities [56]. The currently best-performing method on many benchmarks uses a heatmap-based framework together with optimization of the loss function to foreground versus background pixels [47]. Such supervised methods will typically require large amounts of *labelled* training data in order to generalize across the variability in facial appearance (see [10] for an architecture using high-resolution deep cascades that tries to address this issue).

2.2. Semi-supervised methods

In addition to changes to the network architecture, the issue of lack of training data and inconsistent labeling quality is addressed in semi-supervised models [24, 59] that augment the training process to make use of partially- or weakly-annotated data. Data augmentation based on landmark perturbation [29] or from generating additional views from a 3D face model [60] can be applied to generate more robust pseudo landmark labels. [14] uses constraints from temporal consistency of landmarks based on optic flow to enhance the training of the landmark detector - see also [53]. In [36, 55], multi-task frameworks are proposed in which attribute-networks tasked with predicting other facial attributes including pose and emotion are trained together with the landmark network, allowing for gradient transfer from one network to the other. Similar to this, [35] show improvements using data augmentation with style-translated examples during training. In [15], a teacher-supervises-students (TS³) framework is proposed in which a teacher is trained to filter student-generated landmark pseudolabels into "qualified" and "unqualified" samples, such that the student detectors can retrain themselves with better-quality data. Similarly, in [39], a GAN framework produces "fake" heatmaps that the main branch of the network needs to discriminate, hence improving performance.

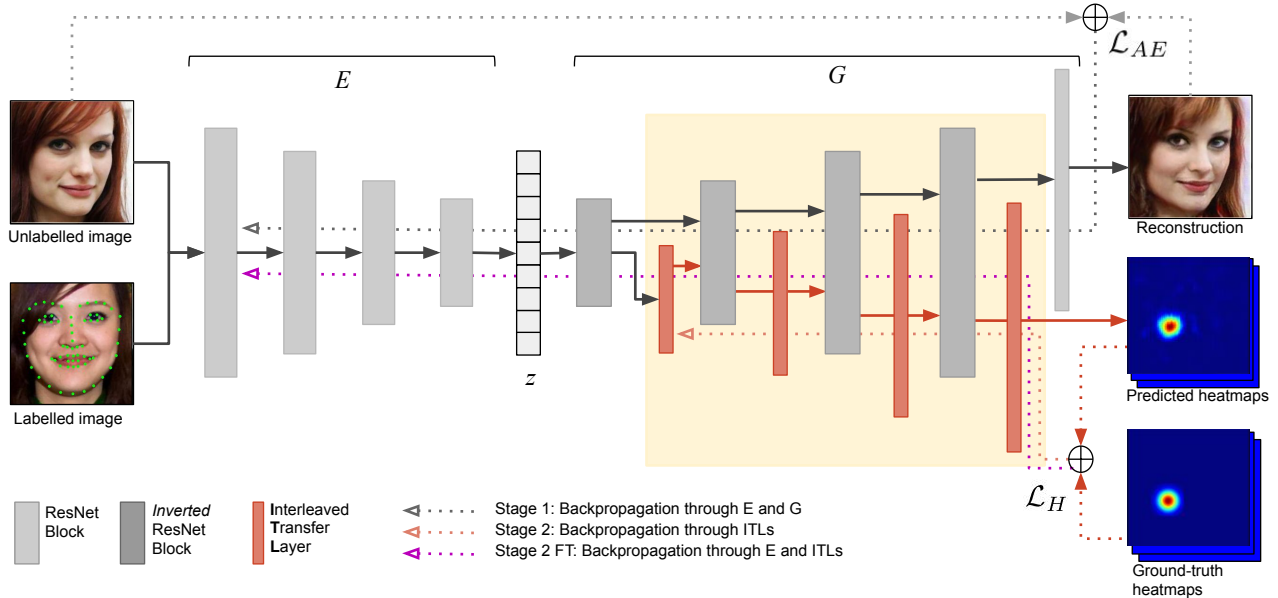


Figure 2: Overview of the 3FabRec pipeline, including the architecture of the autoencoder, as well as training paths for unsupervised, supervised, and the fine-tuning stages (see text for more details).

3. Methods

3.1. Our approach

Most of the semi-supervised approaches discussed above use data augmentation on the same dataset as done for testing. Our approach (see Figs. 1,2) starts from an unsupervised method in which we leverage the *implicit* knowledge about face shape contained in large datasets of faces (such as used for face identification [6]). This knowledge is captured in a low-dimensional latent space of an autoencoder framework. Importantly, the autoencoder also has generative capabilities, i.e., it is tasked during training to reconstruct the face from the corresponding latent vector. This step is done because the following, supervised stage implements a hybrid reconstruction pipeline that uses the generator together with interleaved transfer layers to both reconstruct the face as well as probabilistic landmark heatmaps. Hence, the changes in the latent vector space will be mapped to the position of the landmarks trained on labeled datasets. Given that the first, unsupervised stage has already captured knowledge about facial appearance and face shape, this information will be quickly made explicit during the second, supervised stage allowing for generalization across multiple datasets and enabling low-shot and few-shot training.

3.2. Unsupervised face representation

The unsupervised training step follows the framework of [4] in which an adversarial autoencoder is trained through four loss functions balancing faithful image reconstruction with the generalizability and smoothness of the embedding

space needed for the generation of novel faces. A reconstruction loss \mathcal{L}_{rec} penalizes reconstruction errors through a pixel-based $L1$ error. An encoding feature loss \mathcal{L}_{enc} [19] ensures the creation of a smooth and continuous latent space. An adversarial feature loss \mathcal{L}_{adv} pushes the encoder E and generator G to produce reconstructions with high fidelity since training of generative models using only image reconstruction losses typically leads to blurred images.

As the predicted landmark locations in our method follow directly from the locations of reconstructed facial elements, our main priority in training the autoencoder lies in the accurate reconstruction of such features. Thus, we trade some of the generative power against reconstruction accuracy by replacing the generative image loss, \mathcal{L}_{gen} , used in [4] with a new structural image loss \mathcal{L}_{cs} .

Structural image loss: To penalize reconstructions that do not align facial structures well with input images, we add a structural image loss based on the SSIM [48] image similarity metric, which measures contrast $c(a, b)$ and correlation $s(a, b)$ between two image windows a and b :

$$c(a, b) = \frac{2\sigma_a\sigma_b + c}{\sigma_a^2 + \sigma_b^2 + c}, \quad s(a, b) = \frac{\sigma_{ab} + c/2}{\sigma_a\sigma_b + c/2} \quad (1)$$

The values σ_a and σ_b denote intensity variances of windows a, b and σ_{ab} denotes their covariance. The constant c adds stability against small denominators. It is set to $c = 255^{0.01}$ for images with 8-bit channels. The calculation is run for each $k \times k$ window across the images:

$$cs(x, y) = \frac{1}{|w|} \sum_w c(x_w, y_w) s(x_w, y_w) \quad (2)$$

We obtain the structural image loss by evaluating $cs(x, y)$ with the original image and its reconstructions:

$$\mathcal{L}_{cs}(E, G) = \mathbb{E}_{x \sim p(x)} [cs(x, G(E(x)))] \quad (3)$$

This loss improves the alignment of high-frequency image elements and imposes a penalty for high-frequency noise introduced by the adversarial image loss. Hence, \mathcal{L}_{cs} also serves as a regularizer, stabilizing adversarial training.

Full autoencoder objective: The final training objective is a weighted combination of all loss terms:

$$\begin{aligned} \min_{E, G} \max_{D_z, D_x} \mathcal{L}_{AE}(E, G, D_z, D_x) = \\ \lambda_{rec} \mathcal{L}_{rec}(E, G) + \lambda_{cs} \mathcal{L}_{cs}(E, G) \\ + \lambda_{enc} \mathcal{L}_{enc}(E, D_z) + \lambda_{adv} \mathcal{L}_{adv}(E, G, D_x) \end{aligned} \quad (4)$$

We set λ_{enc} and λ_{adv} to 1.0. λ_{rec} and λ_{cs} are selected so the corresponding loss terms yield similarly large values to each other, while at the same time ensuring a roughly 10 times higher weight in comparison to λ_{enc} and λ_{adv} (given the range of loss terms, we set $\lambda_{rec} \approx 1.0$, $\lambda_{cs} \approx 60.0$).

3.3. Supervised landmark discovery

For landmark detection, we are not primarily interested in producing a RGB image but rather an L -channel image containing landmark probability maps. This can be seen as a form of style transfer in which the appearance of the generated face is converted to a representation that allows us to read off landmark positions. Hence, information about face shape that was implicitly present in the generation of color images before is now made explicit. Our goal is to create this transfer without losing the face knowledge distilled from the very large set of (unlabeled) images as the annotated datasets available for landmark prediction are only a fraction of that size and suffer from imprecise and inconsistent human annotations [14]. For this, we introduce additional, interleaved transfer layers into the generator G .

3.3.1 Interleaved transfer layers

Training of landmark generation starts by freezing all parameters of the autoencoder. We then interleave the inverted ResNet layers of the generator with 3×3 convolutional layers. Each of these Interleaved Transfer Layers (ITL) produces the same number of output channels as the original ResNet layer. Activations produced by a ResNet layer are transformed by these layers and fed into the next higher block. The last convolutional layer mapping to RGB

images is replaced by a convolutional layer mapping to L -channel heatmap images (L = number of landmarks to be predicted). This approach adds just enough flexibility to the generator to produce new heatmap outputs by re-using the pre-trained autoencoder weights.

Given an annotated face image x , the ground truth heatmap H_i for each landmark $l_i \in \mathbb{R}^2$ consists of a 2D Normal distribution centered at l_i and a standard deviation of σ . During landmark training and inference the activations a_1 produced by the first inverted ResNet layer for an encoded image $z = E(x)$ are passed to the first ITL layer. This will transfer the activations and feed it into the next, frozen inverted ResNet layer, such that the full cascade of ResNet and ITLs can reconstruct a landmark heatmap \tilde{H} . The heatmap prediction loss \mathcal{L}_H is defined as the L_2 distance between predicted (\tilde{H}) and ground truth heatmap (H)

$$\mathcal{L}_H(ITL) = \mathbb{E}_{x \sim p(x)} [\|H - ITL(a_1)\|_2] \quad (5)$$

The position of the landmark is $\tilde{l}_i = \underset{u, v}{\operatorname{argmax}} \tilde{H}_i(u, v)$.

3.3.2 Encoder finetuning

Once training of the ITL layers reaches convergence we can perform an optional finetuning step. For this, the encoder E is unfrozen so that ITL layers and encoder are optimized in tandem (see Fig.2).

$$\mathcal{L}_H(ITL) \rightarrow \mathcal{L}_H(E, ITL) \quad (6)$$

Since the updates are only based on landmark errors, this will push E to encode input faces such that facial features are placed more precisely in reconstructed faces. At the same time, other attributes like gender, skin color, or illumination may be removed as these are not relevant for the landmark prediction task. Overfitting is avoided since the generator remains unchanged, which acts as a regularizer and limits the flexibility of the encoder.

4. Experiments¹

4.1. Datasets

VGGFace2 & AffectNet The dataset used for unsupervised training of the generative autoencoder combines two datasets: the VGGFace2 dataset [6], which contains a total of 3.3 million faces collected with large variability in pose, age, illumination, and ethnicity in mind. From the full dataset, we removed faces with a height of less than 100 pixels resulting in 1.8 million faces (from 8631 unique identities). In addition, we add the AffectNet dataset [34] that was designed for capturing a wide variety of facial expressions (thus providing additional variability in face shape), which contains 228k images, yielding a total of 2.1M images for autoencoder training.

¹For experiments on parameter tuning, cross-database results, and further ablation studies, see supplementary materials.

300-W This dataset was assembled by [40] from several sources, including LFPW [2], AFW [26], HELEN [61], XM2VTS [32], and own data and annotated semi-automatically with 68 facial landmarks. Using the established splits reported in [38], a total of 3,148 training images and 689 testing images were used in our experiments. The latter is further split into 554 images that constitute the *common* subset and a further 135 images that constitute the *challenging* subset. Additionally, 300-W contains 300 indoor and 300 outdoor images that define the *private testset* of the original 300-W challenge.

AFLW This dataset [25] contains 24,386 in-the-wild faces with an especially wide range of face poses (yaw angles from -120° – 120°) and roll and pitch angles from -90° – 90°). Following common convention, we used splits of 20,000 images for training and 4,386 for testing and trained with only 19 of the 21 annotated landmarks [28].

WFLW The newest dataset in our evaluation protocol is from [49] containing a total of 10,000 faces with a 7,500/2,500 train/test split. Images were sourced from the WIDER FACE dataset [52] and were manually annotated with a much larger number of 98 landmarks. The dataset contains different (partially overlapping) test subsets for evaluation where each subset varies in pose, expression, illumination, make-up, occlusion, or blur.

4.2. Experimental settings

4.2.1 Unsupervised autoencoder training

Network architecture Our implementation is based on [4] which combines a standard ResNet-18 as encoder with an inverted ResNet-18 (first convolution layers in each block replaced by 4×4 deconvolution layers) as decoder. Both encoder and decoder contain $\approx 10\text{M}$ parameters each. The encoded feature length is 99 dimensions.

Training procedure We train the autoencoder for 50 epochs with an input/output size of 128×128 and a batch-size of 100 images. Upon convergence we add an additional ResNet layer to both the encoder and decoder and train for another 50 epochs with an image size of 256×256 to increase reconstruction fidelity with a batchsize of 50. We use the Adam optimizer [23] ($\beta_1 = 0.0, \beta_2 = 0.999$) with a constant learning rate of 2×10^{-5} , which yielded robust settings for adversarial learning. We apply data augmentations of random horizontal flipping ($p = 0.5$), translation ($\pm 4\%$) resizing (94% to 103%), rotation ($\pm 45^\circ$).

4.2.2 Supervised landmark training

Images are cropped using supplied bounding boxes and resized to 256×256 . For creating ground truth heatmaps, we set $\sigma = 7$. In all experiments we train four ITL layers and generate landmark heatmaps of size 128×128 by



Figure 3: Randomly-generated faces with overlaid generated landmark probability maps.

skipping the last generator layer (as detailed in 4.6, higher generator layers contain mostly decorrelated local appearance information). To train from the landmark dataset images, we apply data augmentations of random horizontal flipping ($p = 0.5$), translation ($\pm 4\%$) resizing ($\pm 5\%$), rotation ($\pm 30^\circ$), and occlusion (at inference time no augmentation is performed). The learning rate during ITL-only training is set to 0.001. During the optional finetuning stage we lower ITL learning rate to 0.0001 while keeping the encoder learning rate the same as during training ($= 2 \times 10^{-5}$) and resetting Adam’s β_1 to the default value of 0.9.

Evaluation Metrics Performance of facial landmark detection is reported here using normalized mean error (NME), failure rate (FR) at 10% NME and area-under-the-curve (AUC) of the Cumulative Error Distribution (CED) curve. For 300-W and WFLW we use the distance between the outer eye-corners as the “inter-ocular” normalization. Due to the high number of profile faces in AFLW, errors are normalized using the width of the (square) bounding boxes following [57].

4.3. Qualitative results

The trained generator is able to produce a wide range of realistic faces from a low-dimensional (99D) latent feature vector z - this is shown in Fig.3 with *randomly-generated* faces with overlaid, predicted landmark heatmaps. To achieve this, the model must have learned inherent information about the underlying structure of faces. We can further illustrate the implicit face shape knowledge by interpolating between face embeddings and observing that facial structures (such as mouth corners) in produced images are constructed in a highly consistent manner (see Fig. 4 for a visualization). This leads to two insights: First, facial structures are actually encoded in the low-dimensional representation z . Second, this information can be transformed into 2D maps of pixel intensities (i.e., a color image) while maintaining high correlation with the originating encoding.

Further examples of the reconstruction quality on challenging images are shown in Fig. 5. As can be seen, the pipeline will try to reconstruct a full face as much as possible given the input, removing occlusions and make-up and




Figure 4: Predicted landmarks of generated faces by interpolation between embedded feature vectors.




Figure 5: 3FabRec results on challenging test examples from WFLW. Rows show the original, and the reconstruction itself, with predicted landmarks, with ground-truth landmarks, and with predicted landmark heatmaps, respectively. The fifth column illustrates a failure case. For more examples, see supplementary materials.

even "upsampling" the face (Fig. 5, first column) in the process. This is because the databases for training the auto-encoder contained mostly unoccluded and non-disguised faces at roughly similar resolutions. Additionally we note that the reconstructed faces will not necessarily preserve the identity as the goal of the fully-trained pipeline is to reconstruct the best-fitting face shape. Although our method is able to handle considerable variations in resolution (Fig. 5, first column), make-up (Fig. 5, second column), lighting (Fig. 5, third column), and pose (Fig. 5, fourth column), it does produce failed predictions in cases when these factors become too extreme, as shown in the fifth column of Fig. 5. Landmark prediction, however, typically degrades gracefully in these cases as the confidence encoded in the heatmaps will also be low.

4.4. Comparison with state-of-the-art

Table 1 shows comparisons of our semi-supervised pipeline with state-of-the-art on the 300-W and the AFLW datasets using the full amount of training. We achieve top-

Method	AFLW		300-W		
	Full	Frontal	Com	Chall.	Full
SDM [51]	4.05	2.94	5.57	15.40	7.52
LBF [37]	4.25	2.74	4.95	11.98	6.32
CFSS [58]	3.92	2.68	4.73	9.98	5.76
Two-Stage [28]	2.17	-	4.36	7.56	4.99
DSRN [33]	1.86	-	4.12	9.68	5.21
SBR [16]	2.14	2.07	3.28	7.58	4.10
SAN [14]	1.91	1.85	3.34	6.60	3.98
LAB [49]	1.85	1.62	2.98	5.19	3.49
ODN [56]	1.63	1.38	3.56	6.67	4.17
LaplaceKL (70K) [39]	1.97	-	<u>3.19</u>	6.87	3.91
3FabRec	<u>1.84</u>	<u>1.59</u>	3.36	<u>5.74</u>	<u>3.82</u>

Table 1: Normalized mean error (%) on 300-W dataset. Best results highlighted in bold, second best are underlined.

Method	AUC	FR
M ³ CSR [11]	47.52	5.5
CFSS [57]	49.87	5.05
DenseReg+MDM [1]	52.19	3.67
JMFA [12]	54.85	1.00
LAB [49]	58.85	0.83
3FabRec	54.61	0.17

Table 2: Area under the curve (AUC) and failure rate (FR in (%) @0.1) on the 300-W testset.

2 accuracy on nearly all test sets with the exception of the common set from 300-W. This demonstrates that our framework is able to reach current levels of performance despite a much lighter, supervised training stage using only a few interleaved transfer layers on top of the generator pipeline.

The results in Table 2 for AUC and FR for the commonly-reported 300-W dataset demonstrate that our framework achieves the lowest failure rate of all methods (our FR=0.17 corresponds to only 1 image out of the full set that has large enough errors to count as a failure). At the same time, the AUC is in the upper range but not quite as good as that of [49], for example, which means that overall errors across landmarks are low, but more equally distributed compared to the top-performing methods.

The NME results in Table 3 show that on the newest WFLW dataset, our approach performs at levels of the LAB method [49] with most subsets, although we perform consistently below the current StyleAlign approach (SA, [35]) - note, however, that this approach could be easily implemented into our framework as well, which would allow us to disentangle the 99D-feature vector into style attributes [4]

	Method	Full	Pose	Exp.	Ill.	Mk.	Occ.	Blur
						Up		
NME (%)	SDM [51]	10.29	24.10	11.45	9.32	9.38	13.03	11.28
	CFSS [58]	9.07	21.36	10.09	8.30	8.74	11.76	9.96
	DVLN [50]	6.08	11.54	6.78	5.73	5.98	7.33	6.88
	LAB [49]	5.27	10.24	5.51	5.23	5.15	6.79	6.32
	SAN [14]	5.22	10.39	5.71	5.19	5.49	6.83	5.80
	Wing [49]	5.11	8.75	5.36	4.93	5.41	6.37	5.81
	SA [35]	4.39	8.24	4.68	4.24	4.37	5.60	4.86
3FabRec	5.62	10.23	6.09	5.55	5.68	6.92	6.38	
FR @0.1 (%)	SDM [51]	29.40	84.36	33.44	26.22	27.67	41.85	35.32
	CFSS [58]	20.56	66.26	23.25	17.34	21.84	32.88	23.67
	DVLN [50]	10.84	46.93	11.15	7.31	11.65	16.30	13.71
	LAB [49]	7.56	28.83	6.37	6.73	7.77	13.72	10.74
	SAN [14]	6.32	27.91	7.01	4.87	6.31	11.28	6.60
	Wing [49]	6.00	22.70	4.78	4.30	7.77	12.50	7.76
	SA [35]	4.08	18.10	4.46	2.72	4.37	7.74	4.40
3FabRec	8.28	34.35	8.28	6.73	10.19	15.08	9.44	
AUC @0.1	SDM [51]	0.300	0.023	0.229	0.324	0.312	0.206	0.239
	CFSS [58]	0.366	0.063	0.316	0.385	0.369	0.269	0.304
	DVLN [50]	0.455	0.147	0.389	0.474	0.449	0.379	0.397
	LAB [49]	0.532	0.235	0.495	0.543	0.539	0.449	0.463
	SAN [15]	0.536	0.236	0.462	0.555	0.522	0.456	0.493
	Wing [49]	0.534	0.310	0.496	0.541	0.558	0.489	0.492
	SA [35]	0.591	0.311	0.549	0.609	0.581	0.516	0.551
3FabRec	0.484	0.192	0.448	0.496	0.473	0.398	0.434	

Table 3: Evaluation results on WFLW dataset.

to generate augmented training data). The main reason for this is that WFLW contains much more heavy occlusions and extreme appearance changes compared to our training sets leading to more failure cases (see Fig.5 fifth column).

4.5. Limited training data and few-shot learning

Tables 4, 5, 6 showcase the central result of our framework: when training on only parts of the training set, 3FabRec can beat the published benchmark performance values.

300-W Table 4 shows that performance is comparable to that of 2-year-old approaches trained on the full dataset (cf. Table 1) although 3FabRec was trained only with 10% of the dataset. In addition, performance does not decrease much when going to lower values of 5% and 1.5% of training set size. Even when training with only 10 images or 1 image, our approach is able to deliver reasonably robust results (see Fig.1 for landmark reconstruction results from training with 10 images).

AFLW For this dataset (Table 5), our approach already starts to come ahead at 20% of training set size with little degradation down to 1%. Again, even with only a few images 3FabRec can make landmark predictions.

WFLW For this more challenging dataset (Table 6), our approach easily outperforms the StyleAlign [21] method as soon as less than 10% is used for training while being able




Figure 6: Layer-analysis of 3FabRec. Gray curve: cumulative number of network parameters; blue curve: spatial dimension of each layer. The four red blocks indicate the ITL layers, with arrows showing how well the landmark heatmap can be predicted when starting from that layer.

to maintain landmark prediction capabilities down to only 10 images in the training set.

4.6. Ablation studies

4.6.1 Effects of ITLs

In order to see where information about landmarks is learned in the interleaved transfer layers, Figure 6 shows the reconstruction of the landmark heatmap when using all four layers versus decreasing subsets of the upper layers. As can be seen, the highest layer has only very localized information (mostly centered on eyes and mouth), whereas the lower layers are able to add information about the outlines - especially below layer 2.

Localization accuracy is reported on the 300-W dataset (NME of 51 inner landmarks and outlines, as well as FR) in Table 7. As can be expected from the visualization, performance is bad for the upper layers only, but quickly recovers (especially when including the outlines) below layer 2. The reason for this is that the upper layers of the generator will mostly contain localized, de-correlated information at the pixel level, whereas the lower layers are closer to the more global and contextual information necessary to cover highly variable outlines (cf. blue curve in Figure 6, note that all ITLs have 3×3 convolutions). As the gray curve in Figure 6 and Table 7 show as well, the ITLs can achieve this with only very few additional parameters.

4.6.2 Effects of finetuning

Table 8 reports the effects of running the model with and without finetuning on the full testsets of the three evaluated datasets. The additional retraining of the autoencoder allows for better reconstruction of the faces and results in benefits of 10.9% on average (8.9% for 300-W, 15.2% for AFLW, and 8.5% for WFLW, respectively).

300-W dataset																					
Method	Training set size																				
	100%			20%			10%			5%			50 (1.5%)			10 (0.3%)			1 (0.003%)		
RCN+ [21]	4.20	7.78	4.90	-	9.56	5.88	-	10.35	6.32	-	15.54	7.22	-	-	-	-	-	-	-	-	-
RCN+ [21] [†]	3.00	4.98	3.46	-	6.12	4.15	-	6.63	4.47	-	9.95	5.11	-	-	-	-	-	-	-	-	-
SA [35]	3.21	6.49	3.86	3.85	-	-	4.27	-	-	6.32	-	-	-	-	-	-	-	-	-	-	-
TS ³ [15]	2.91	5.9	3.49	4.31	7.97	5.03	4.67	9.26	5.64	-	-	-	-	-	-	-	-	-	-	-	-
3FabRec	3.36	5.74	3.82	3.76	6.53	4.31	3.88	6.88	4.47	4.22	6.95	4.75	4.55	7.39	5.10	4.96	8.29	5.61	8.45	15.84	9.92

Table 4: NME (%) with reduced training sets on 300-W. [†] RCN+ reports errors normalized by eye-center distance - for better comparison values were rescaled by the known ratios of inter-ocular to inter-pupil distances, "-" denotes values not reported.

AFLW dataset																
Method	Training set size															
	100%		20%		10%		5%		1%		50 (0.0025%)		10 (0.0005%)		1 (<0.0001%)	
RCN+ [21]	1.61	-	-	-	-	-	2.17	-	2.88	-	-	-	-	-	-	
TS ³ [15]	-	-	1.99	1.86	2.14	1.94	2.19	2.03	-	-	-	-	-	-	-	
3FabRec	1.87	1.59	1.96	1.74	2.03	1.74	2.13	1.86	2.38	2.03	2.74	2.23	3.05	2.56	4.93	4.04

Table 5: NME (%) with reduced training sets for AFLW. The first column in each cell is the full testset, the second is the frontal testset, "-" denotes values not reported.

WFLW dataset							
Method	Training set size						
	100%	20%	10%	5%	50	10	1
SA [21]	4.39	6.00	7.20	-	-	-	-
3FabRec	5.62	6.51	6.73	7.68	8.39	9.66	15.79

Table 6: NME (%) with reduced training sets for WFLW.

	Trained ITLs			
	1+2+3+4	2+3+4	3+4	4
Input size	256x8x8	128x16x16	64x32x32	64x64x64
Trainable params	881k	291k	143k	106k
300-W NME $\neg O$	3.54	3.63	5.34	16.34
300-W NME O	6.58	7.32	18.17	40.24
300-W FR@0.1	1.45	2.03	22.93	91.44

Table 7: Parameters and training results for ITLs ($\neg O$ =without outlines, O =outlines only)

	300-W	AFLW	WFLW
NME before FT	4.16	2.12	6.11
NME after FT	3.82	1.84	5.62

Table 8: NME (%) before and after finetuning on full testsets.

4.7. Runtime performance

Since inference complexity is equivalent to two forward-passes through a ResNet-18, our method is able to run at frame rates of close to 300fps on a TitanX GPU - an order of magnitude faster than state-of-the-art approaches with similar, high accuracy (LAB [49]=16fps, Wing [18]=30fps, Deep Regression [28]=83fps, Laplace [39]=20fps).

5. Conclusion

With 3FabRec, we have demonstrated that an unsupervised, generative training on large amounts of faces captures implicit information about face shape, making it possible to solve landmark localization with only a minimal amount of supervised follow-up training. This paradigm makes our approach inherently more robust against overfitting to specific training datasets as well as against human annotation variability [14]. The critical ingredients of 3FabRec that enable this generalization are the use of an adversarial autoencoder that reconstructs high-quality faces from a low-dimensional latent space, together with low-overhead, interleaved transfer layers added to the generator stage that transfer face reconstruction to landmark heatmap reconstruction.

Results show that the autoencoder is easily able to generalize from its unlabeled training set to data from unseen datasets. This offers generalization for training from only a few percent of the training set and still produces reliable results from only a few annotated images - far below anything reported so far in the literature. At the same time, since inference amounts to only two forward passes through a ResNet18, our method achieves much higher runtime performance than other highly accurate methods.

Acknowledgements This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. 2019-0-00079, Department of Artificial Intelligence, Korea University)

References

- [1] Riza Alp Guler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, and Iasonas Kokkinos. Densereg: Fully convolutional dense shape regression in-the-wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6799–6808, 2017. 6
- [2] Peter N Belhumeur, David W Jacobs, David J Kriegman, and Neeraj Kumar. Localizing parts of faces using a consensus of exemplars. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2930–2940, 2013. 5
- [3] Matteo Bordini. A review of facial landmark extraction in 2d images and videos using deep learning. *Big Data and Cognitive Computing*, 3(1):14, 2019. 1
- [4] Björn Browatzki and Christian Wallraven. Robust discrimination and generation of faces using compact, disentangled embeddings. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3, 5, 6
- [5] Adrian Bulat and Georgios Tzimiropoulos. Two-stage convolutional part heatmap regression for the 1st 3D face alignment in the wild (3DFAW) challenge. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9914 LNCS, pages 616–624, 2016. 2
- [6] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *Face & Gesture Recognition (FG 2018)*, pages 67–74. IEEE, 2018. 3, 4
- [7] Gavin C Cawley and Nicola LC Talbot. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(Jul):2079–2107, 2010. 2
- [8] Timothy F Cootes, Gareth J Edwards, and Christopher J Taylor. Active appearance models. In *European Conference on Computer Vision*, pages 484–498. Springer, 1998. 2
- [9] Timothy F Cootes and Christopher J Taylor. Active shape modelssmart snakes. In *BMVC92*, pages 266–275. Springer, 1992. 2
- [10] Arnaud Dapogny, Kévin Bailly, and Matthieu Cord. DeCaFA: Deep Convolutional Cascade for Face Alignment In The Wild. pages 6893–6901, 2019. 2
- [11] Jiankang Deng, Qingshan Liu, Jing Yang, and Dacheng Tao. M3 csr: Multi-view, multi-scale and multi-component cascade shape regression. *Image and Vision Computing*, 47:19–26, 2016. 6
- [12] Jiankang Deng, George Trigeorgis, Yuxiang Zhou, and Stefanos Zafeiriou. Joint multi-view face alignment in the wild. *IEEE Transactions on Image Processing*, 28(7):3636–3648, 2019. 2, 6
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [14] Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Style Aggregated Network for Facial Landmark Detection. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 379–388, 2018. 2, 4, 6, 7, 8
- [15] Xuanyi Dong and Yi Yang. Teacher Supervises Students How to Learn From Partially Labeled Images for Facial Landmark Detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 783–792, 2019. 2, 7, 8
- [16] Xuanyi Dong, Shouou-I Yu, Xinshuo Weng, Shih-En Wei, Yi Yang, and Yaser Sheikh. Supervision-by-registration: An unsupervised approach to improve the precision of facial landmark detectors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 360–368, 2018. 6
- [17] Zhen-Hua Feng, Guosheng Hu, Josef Kittler, William Christmas, and Xiao-Jun Wu. Cascaded collaborative regression for robust facial landmark detection trained using a mixture of synthetic and real images with dynamic weighting. *IEEE Transactions on Image Processing*, 24(11):3425–3440, 2015. 2
- [18] Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, and Xiao-Jun Wu. Wing Loss for Robust Facial Landmark Localisation with Convolutional Neural Networks. 2017. 8
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014. 3
- [20] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7):1527–1554, 2006. 2
- [21] Sina Honari, Pavlo Molchanov, Stephen Tyree, Pascal Vincent, Christopher Pal, and Jan Kautz. Improving Landmark Localization with Semi-Supervised Learning. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1546–1555, 2018. 7, 8
- [22] Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*, 2018. 2
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [24] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pages 3581–3589, 2014. 2
- [25] Martin Koestinger, Paul Wohlhart, Peter M Roth, and Horst Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV workshops)*, pages 2144–2151. IEEE, 2011. 5
- [26] Vuong Le, Jonathan Brandt, Zhe Lin, Lubomir Bourdev, and Thomas S Huang. Interactive facial feature localization. In *European Conference on Computer Vision*, pages 679–692. Springer, 2012. 5

- [27] Zhujin Liang, Shengyong Ding, and Liang Lin. Unconstrained facial landmark localization with backbone-branches fully-convolutional networks. *arXiv preprint arXiv:1507.03409*, 2015. [2](#)
- [28] Jiangjing Lv, Xiaohu Shao, Junliang Xing, Cheng Cheng, and Xi Zhou. A deep regression architecture with two-stage re-initialization for high performance facial landmark detection. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-January:3691–3700, 2017. [2](#), [5](#), [6](#), [8](#)
- [29] Jiang-Jing Lv, Cheng Cheng, Guo-Dong Tian, Xiang-Dong Zhou, and Xi Zhou. Landmark perturbation-based data augmentation for unconstrained face recognition. *Signal Processing: Image Communication*, 47:465–475, 2016. [2](#)
- [30] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015. [2](#)
- [31] Daniel Merget, Matthias Rock, and Gerhard Rigoll. Robust Facial Landmark Detection via a Fully-Convolutional Local-Global Context Network. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 781–790, 2018. [2](#)
- [32] Kieron Messer, Jiri Matas, Josef Kittler, Juergen Luetttin, and Gilbert Maitre. Xm2vtsdb: The extended m2vts database. In *Second international conference on audio and video-based biometric person authentication*, volume 964, pages 965–966, 1999. [5](#)
- [33] Xin Miao, Xiantong Zhen, Xianglong Liu, Cheng Deng, Vasilis Athitsos, and Heng Huang. Direct shape regression networks for end-to-end face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5040–5049, 2018. [6](#)
- [34] Ali Mollahosseini, Behzad Hasani, and Mohammad H. Mahoor. AffectNet: A Database for Facial Expression, Valence, and Arousal Computing in the Wild. *IEEE Transactions on Affective Computing*, 2017. [4](#)
- [35] Shengju Qian, Keqiang Sun, Wayne Wu, Chen Qian, and Ji-aya Jia. Aggregation via Separation: Boosting Facial Landmark Detector with Semi-Supervised Style Translation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. [2](#), [6](#), [7](#), [8](#)
- [36] Rajeev Ranjan, Vishal M Patel, and Rama Chellappa. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(1):121–135, 2017. [2](#)
- [37] Shaoqing Ren, Xudong Cao, Yichen Wei, and Jian Sun. Face alignment at 3000 fps via regressing local binary features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1685–1692, 2014. [6](#)
- [38] Shaoqing Ren, Xudong Cao, Yichen Wei, and Jian Sun. Face alignment via regressing local binary features. *IEEE Transactions on Image Processing*, 25(3):1233–1245, 2016. [5](#)
- [39] Joseph P Robinson, Yuncheng Li, Ning Zhang, Yun Fu, and Sergey Tulyakov. Laplace Landmark Localization. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. [2](#), [6](#), [8](#)
- [40] Christos Sagonas, Epameinondas Antonakos, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 Faces In-The-Wild Challenge: database and results. *Image and Vision Computing*, 47:3–18, 2016. [5](#)
- [41] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 397–403, 2013. [1](#)
- [42] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep convolutional network cascade for facial point detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3476–3483, 2013. [2](#)
- [43] James Thewlis, Samuel Albanie, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of landmarks by descriptor vector exchange. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6361–6371, 2019. [2](#)
- [44] Phil Tresadern, Tim Cootes, Chris Taylor, and Vladimir Petrović. Face alignment models. In *Handbook of face recognition*, pages 109–135. Springer, 2011. [1](#)
- [45] George Trigeorgis, Patrick Snape, Mihalis A Nicolaou, Epameinondas Antonakos, and Stefanos Zafeiriou. Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4177–4187, 2016. [2](#)
- [46] Tim Valentine, Michael B Lewis, and Peter J Hills. Face-space: A unifying concept in face recognition research. *The Quarterly Journal of Experimental Psychology*, 69(10):1996–2019, 2016. [2](#)
- [47] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive Wing Loss for Robust Face Alignment via Heatmap Regression. 2019. [2](#)
- [48] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. IEEE, 2003. [3](#)
- [49] Wenyang Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, and Qiang Zhou. Look at Boundary: A Boundary-Aware Face Alignment Algorithm. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2129–2138, 2018. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [50] Yue Wu, Tal Hassner, KangGeon Kim, Gerard Medioni, and Prem Natarajan. Facial Landmark Detection with Tweaked Convolutional Neural Networks. 2015. [7](#)
- [51] Xuehan Xiong and Fernando De la Torre. Supervised descent method and its applications to face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 532–539, 2013. [1](#), [2](#), [6](#), [7](#)
- [52] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1281–1290, 2017. [5](#)
- [53] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-Shot Adversarial Learning of Realistic Neural Talking Head Models. 2019. [2](#)

- [54] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1058–1067, 2017. [2](#)
- [55] Zhanpeng Zhang, Ping Luo, Chen-Change Loy, and Xiaoou Tang. Facial Landmark Detection by Deep Multi-task Learning. *Eccv*, 2014. [2](#)
- [56] Meilu Zhu, Daming Shi, Muhammad Sadiq, and Mingjie Zheng. Robust Facial Landmark Detection via Occlusion-adaptive Deep Networks. *Cvpr*, pages 3486–3496, 2019. [2](#), [6](#)
- [57] Shizhan Zhu, Cheng Li, Chen Change Loy, and Xiaoou Tang. Face alignment by coarse-to-fine shape searching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4998–5006, 2015. [5](#), [6](#)
- [58] Shizhan Zhu, Cheng Li, Chen Change Loy, and Xiaoou Tang. Face alignment by coarse-to-fine shape searching. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 07-12-June:4998–5006, 2015. [6](#), [7](#)
- [59] Xiaojin Zhu. Semi-supervised learning. *Encyclopedia of Machine Learning and Data Mining*, pages 1142–1147, 2017. [2](#)
- [60] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z. Li. Face Alignment Across Large Poses: A 3D Solution. 2016. [2](#)
- [61] Xiangxin Zhu and Deva Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2879–2886. IEEE, 2012. [5](#)